

DEVICE AND METHOD FOR DATA TIMESTAMPING

BACKGROUND OF THE INVENTION

5

1. Field of the Invention

This invention relates to a device adapted to provide data time-stamping and a method for providing data time-stamping. More particularly, but not
10 exclusively, it relates to a device and method for providing time-stamping without recourse to a trusted third party.

It will be appreciated that any references to data or data set herein relate to amongst other things, but not exclusively, files, data, documents, and
15 software applications.

2. Description of the prior art

Digital time-stamping is a method whereby an element of data, or data set,
20 can be bound to a particular point in time. To minimise the risk that either the data or the time-stamp can be tampered with at a later date a cryptographic digital signature is used to protect both elements. This is clearly of importance when it is important to provide non-repudiable proof of the existence of data, for example in legal matters such as the formation
25 and agreement of a contract or the timing of a revision of a clause of a contract, or of a will. These are just some examples.

Current time-stamping techniques include a method which relies upon the passing of the data to be time-stamped over a network, such as the Internet,
30 to a trusted timeserver incorporating a trusted clock maintained by a trusted third party, as shown in Figure 1, which time-stamps and digitally signs the data, and sends it back to the originator.

This has security disadvantages in that it involves the transfer over a network, typically the Internet, of the data or time-stamped data which can be intercepted. The data may be altered, re-hashed and sent for time-stamping by the interceptor, thus presenting to a recipient a differently time-stamped data set and associated hash-created digest, which will look correct to the recipient.

Additionally there is the problem of confidence in the trusted third party maintaining the trusted clock. The trusted third party may be certified by an independent Certification Authority. Whilst this gives a high degree of confidence to users, there is a risk that the certificate may be rescinded, expire or be compromised without the immediate knowledge of the users of the trusted data. It will be appreciated that the confidence in the veracity of the timestamp comes from the reputation of the party running the trusted clock and the security of the cryptographic techniques used to sign the hash-created digest.

Remote trusted third party clocks also have a problem of latency (delay) in that a significant amount of time may elapse between the production of data and its time-stamping, it is not an immediate process. There are also limits on throughput in remote trusted third party clocks which can exacerbate the latency problem if the trusted clock forms a constriction in the data flow.

Time-stamping of data by using an internal clock of a computer from which the data originates is generally held to be unacceptable as the internal clock of such computers, such as PC's can be easily altered by simple software alterations.

GENERAL DESCRIPTION OF THE INVENTION

It is an aim of the present invention to provide a data time-stamping device which ameliorates, at least in part, at least one of the above-mentioned
5 disadvantages or problems.

It is another aim of the present invention to provide a method of data time-stamping which ameliorates, at least in part, at least one of the above-mentioned disadvantages or problems.

10

According to a first aspect of the present invention there is provided a storage device including a trusted clock, a memory (or storage media), a time-stamper and a digital signer arranged such that the device is adapted to store to the memory data that has been time-stamped by the time-
15 stamper, with a time obtained from the trusted clock, and digitally signed with a digital signature by the digital signer.

It will be understood that the term "trusted clock" relates to a clock, which is believed to be trustworthy, for example a sealed or otherwise tamper-
20 proof clock unit which is physically and logically difficult or impossible to tamper with, or for example a clock which has its time-stamp authenticity certified by a Certification Authority (CA).

It will also be understood that "data storage device" includes a stand alone
25 device, a sub-system, appliance, system, or local distributed memory network, but does not include internet-distributed memory storage.

The digital signature may be encrypted using asymmetrical encryption, for example PKI, or symmetric encryption, for example DES.

30

- The memory will typically be a long term storage medium, not for example a communication channel (e.g. a data bus) or volatile memory e.g. RAM or a temporary buffer. Long term storage media may include, in a non-exhaustive list, CD, DVD, tape, ZipTM disc, magnetic-optical disc, magnetic disc or any recordable solid state memory such as EPROM, Flash, MRAM, EEPROM or solid state device. The memory, or storage media, may be removable from the storage device or alternatively it may be fixed to/within the storage device.
- 10 The storage device, apparatus, or system could be a simple storage device such as disc drive or tape drive, or a more complex system such as a disc array, disc sub-system, tape library or optical jukebox; or a disaggregated storage network, a storage area network, or a network attached storage device.
- 15 The storage device, apparatus, or system may provide essentially just a storage function, and will in general have no general computational ability or purpose. It will not, for example, be part of the memory of a general purpose server or computer (e.g. not a PC's memory).
- 20 There may be a controller associated with the trusted clock. The controller may have controller logic running thereupon. There may be means of checking the veracity of the controller logic The controller logic may be time-stamped. The controller logic may be time-stamped prior to passing data through the trusted clock. The controller logic signature may be checked prior to the time-stamping of data. This prevents the downloading of fake control data into the controller (known as spoofing) thereby preventing alteration of the clock time.
- 25
- 30 The trusted clock may be mounted upon a plug-in card. The card may be a PCI card. Alternatively the trusted clock may be in the form of a read only device. The clock may have no externally modifiable logic. It may have

essentially only an output time signal. A recalibration input, as possibly the only input signal to the clock, is optional.

The data may or may not be encrypted prior to time-stamping. The encryption could take place within the storage device or externally of the device or system prior to time-stamping by the trusted device (clock).

The system may time-stamp all data that it receives for storage. Alternatively the system may include logic that will apply the use of the time-stamping methodology to selected elements of the data being time-stamped. There may be a flag which indicates that an element of data is to be time-stamped. This flag may be: 1) embedded within the data itself; 2) provided via the command language used for communication between the storage system or device and a host computer (e.g. a SCSI or filter channel command); or 3) provided via a configuration setting of the storage device or system (e.g. a setting on the controller may be turned to and from "time-stamp" and "do not time-stamp").

An output of the time-stamper may be a printer thereby producing a non-alterable, physically secure record of the data, or digest, timestamp and signature.

According to a second aspect of the present invention there is provided a method of storing secure time-stamped data on a data storage device comprising the steps of:

- (i) providing a trusted clock at the data storage device;
- (ii) time-stamping the data at the data storage device;
- (iii) creating a digital signature dependent upon the content of the data and the timestamp; and

(iv) storing the data and associated signature on a recording medium of the data storage device.

5 The digital signature may be encrypted using asymmetric or symmetric encryption. The recording medium may include, in a non-exhaustive list, CD, DVD, Zip TM disc, magnetic-optical disc, magnetic disc or any form of recordable solid state memory such as EPROM, Flash, MRAM, or solid state disc. The storage device, apparatus, or system could be a simple
10 storage device such as disc drive or tape device or a more complex system such as a disc array, disc subsystem, tape library or optical jukebox; or a disaggregated storage network, a storage area network, or network attached storage device. The medium may be removable from the storage device or alternatively may be fixed to/within the storage device.

15 The trusted clock may be provided upon a plug-in card. The card may be a PCI card. Alternatively the trusted clock may be in the form of a read only device.

20 The data may or may not be encrypted prior to time-stamping, and the data plus time stamp is generally cryptographically signed.

According to a third aspect of the present invention there is provided a data storage device or system adapted to time-stamp and store data that it
25 receives, the device being connected to a private or public network, and the device being adapted to receive data from a remote source connected to the network and to time-stamp the data and to store the time-stamped data locally at the data storage device or system without transmitting time-stamped data across the network.

30

Preferably the network may have a plurality of data storage device on it, and at least one of the data storage devices is adapted to time-stamp and store data.

- 5 According to a fourth aspect the invention comprises a method of time-stamping and storing data over a public or private network, the method comprising transmitting data to a data storage device attached to the network and time-stamping the data using a trusted clock and storing the time-stamped data at the data storage device without transmitting time-stamped data across the network.

- 10 According to a fifth aspect of the present invention, there is provided software, firmware, or a computer readable medium having a program recorded thereupon, which, in use, causes a processor of a data storage device running the program to execute a process in accordance with the second aspect of the present invention; or which when operating in a control processor of a data storage device causes that device to be a device in accordance with the first aspect of the invention; or which when running on a data storage device or system that is network-attached causes the method of the fourth aspect of the invention to be performed, or a network in accordance with the third aspect of the invention to be created.

- 20 According to a sixth aspect of the present invention there is provided a data storage device including a trusted clock, the storage device being adapted to store to memory data which has been time-stamped by the clock and which has been digitally signed.

The data storage device may also digitally sign the time-stamped data.

- 30 According to a seventh aspect of the present invention there is provided a method of storing time-stamped data on a network comprising transmitting the data from a first device to a data storage device in accordance with the

first aspect of the present invention and time-stamping and recording the data at the data storage device in the absence of transmitting the time-stamped data back to the first device for storage.

- 5 The invention may have any one or more of the advantages of (i) improving security, i.e. reducing the likelihood of manipulation of the data and timestamp by a third party; (ii) making the time-stamping of data almost instantaneous thereby reducing delays; and (iii) reducing or obviating network bandwidth constraints, increasing throughput of data when
- 10 compared to the prior art arrangements. The prior art arrangements typically have a trusted clock at a point of a network and other network elements, remote from the clock, transmit their data over the network to the trusted clock where it is time-stamped, signed and transmitted back to its originating network element. The present invention further minimises the
- 15 bulk movement of data over a network by having time-stamping at the site where data may be stored. Furthermore, there is a reduced chance of the telecommunications link between the data-originating device and the time-stamped data storage device being interrupted if the time-stamped data is stored at or close to where it is time-stamped. This improves connection
- 20 reliability issues. On congested networks avoiding a "return" transmission leg for the time-stamped data can help avoid loss of packets and can help to reduce congestion.

It will be appreciated that time-stamping can refer to stamping data with a

25 date. It need not, but may, give time in hours, minutes, seconds or subdivisions thereof.

BRIEF DESCRIPTION OF THE DRAWINGS

- 30 The invention will now be described, by way of example, with reference to the accompanying drawings, in which:

Figure 1 is a schematic diagram of a prior art remote trusted third party time-stamping device;

Figure 2 is a schematic diagram of a prior art digital signature scheme;

Figure 3 is a schematic representation of a data time-stamping arrangement according to the present invention;

Figure 4 is a flow diagram showing a data time-stamping method according to the present invention;

Figure 5 is a schematic diagram showing a network with storage devices attached thereto; and

Figure 6 shows another embodiment of the invention.

DESCRIPTION OF THE PREFERRED EMBODIMENT

Current trusted third party time-stamping systems, as shown in Figure 1, involve the transmittal of data over a network to the trusted third party for time-stamping. Data, or a digest of the data, is sent from a computer (e.g. a PC 1) via telecommunications 2 to a network, e.g. the internet 3. The data is routed on the internet 3 to a trusted clock 4 attached to the internet via telecommunications 5 and is time-stamped. Once time-stamped the data may be passed back to the internet via telecommunications 6 and may then be sent via telecommunications 7 to a storage device 8 for storage or it may be sent back to the originator of the data via telecommunications 9 for storage. This introduces delays, has a throughput which is limited by the bandwidth of the network and has opportunities for data interception, connections failures, and falsification of time-stamps.

Digital signatures, see for example Figure 2, reduce the opportunities for data tampering and falsification. This involves passing the data through a hashing algorithm to obtain a digest of the message. A specific digest is almost impossible/very difficult to produce from data other than the original data hashed. The digest is then encrypted using an asymmetric encryption private key to provide a signature. The signature is appended to the data and transmitted with it.

A third party who has the public key which is complementary to the private key used in the encryption process can decrypt the signature to obtain the digest. The third party can rehash the received data and calculate the digest of this. The digest from the signatures and the rehashed digest are compared, if they do not match then the data has been tampered with.

In one embodiment of the present invention, shown in Figure 3, data from data source 10 is passed into a storage device 12. The storage device 12 (with its boundary shown as 13) comprises an interface 14, a data buffer 16, a secure controller 18 with an associated trusted clock/signature module 20, and data storage media 22, 22b, 22c.

The data from the external data source 10 may or may not be encrypted prior to being passed into the storage device 12. The external data source 10 may be for example a LAN, the Internet, a PC or a server.

The interface 14 serves to ensure interoperability and consistent data handling between different data sources 10 and the storage device 12. The interface 14 may take the form of, for example, an internal bus, SCSI or FiberChannel interface. The SCSI commands may have bespoke data control protocols written into them in order to identify data, data types or data sets which require time-stamping.

The data buffer 16 maintains a steady and consistent data transfer rate to the controller 18. The buffer 16 is typically a piece of memory.

5 The secure controller 18 controls the formatting and preparation of data prior to their recording on the media 22a, 22b, 22c. This can include blocking and compression of the data.

10 The data passed to the controller 18 will typically have a flag set which identifies it as requiring time-stamping or not. The controller 18 then either filters out data flagged "time-stamp me" and passes only (or substantially only) the data with the flag set to 'timestamp' to the trusted clock module 20 for time-stamping, or it sends all of the data to the trusted clock which only time-stamps flagged data.

15 The controller 18 may also control the trusted clock 20. Control logic for the controller 18 may be protected by a separate trust mechanism. This may allow the veracity and/or origin of the logic to be checked and may aid in the detection of downloaded fake control logic.

20 The trusted clock module 20 timestamps and digitally signs the data in a conventional manner, for example using DSA, and passes the data back to the controller 20, along with the signature. As will be appreciated, the data could be a digest or signature of a larger set of data. The controller 18 contains a checking routine to confirm that the time-stamping is successful.

25 If it is not correctly time-stamped the data is passed back to the trusted clock module 20 for retime-stamping. The controller 18 writes the data , timestamp and signature to storage media 22a, 22b, 22c, either in a single block or in a fragmented form. If it is written in a fragmented form, there must be data control logic provided in order to locate the fragments.

30

A public key 24 which, corresponds to the private key used in the digital signing of the data is placed on a network 26. A recipient of the data can

obtain the public key 24 from the network 26 or it can be sent to them either via E-mail or on media.

- 5 It will be appreciated that the public key need not be 'published' but may be retained by the author of the data for their own use, or disseminated to a restricted group of people/entities.

- 10 The trusted clock module 20 is typically hardwired into the storage device 12 in order to reduce the likelihood of tampering and bogus insertions of clocks into devices. The clock module 20 may be made tamperproof and/or tamper evident by any convenient method (for example it may be encased in resin or other suitable material to prevent/indicate attempts to access it physically). It is recommended that the trusted clock 20 is certified by a trusted CA, but this is not essential. Other ways of having a trusted clock
15 exist (e.g. an encapsulated clock which cannot be altered and can only output the date and time).

- 20 Provision may made for the replacement of the trusted clock 20 at the expiry of the certificate (e.g. or plug in/out clock module), or authorised service personnel may be capable of removing an encapsulated hardwired clock and replacing it with another, possibly requiring security access codes to disable anti-forgery protection logic. Alternatively it may be possible to upload a new certificate into the clock.

- 25 Provision may be made for the correction of drift of the trusted clock. For example, the clock may be arranged to synchronise itself with a trusted time signal periodically (e.g. with a satellite clock signal).

- 30 An alternative to the hardwiring of the clock module 20 is the use of a removable clock module, for example an insertable plug in – plug out cards containing the clock module. This increases the risk of tampering but has

the advantage of ease of maintenance and replaceability upon the expiry of a certificate period for a particular clock module.

The storage device 13 may be a disc drive, or a tape drive, having no general purpose computing ability, and not being programmable for tasks other than storing and/or retrieving data (with time-stamping and possibly signature generation facilities). Alternatively, whilst still having functionality limited to being essentially a data storage device, it may be more complex such as an array of linked memory stores.

Figure 4 is a flow diagram of a method of time-stamping of data.

Data enters a storage device (Step 50) and is passed to the controller (Step 52). The controller examines the data to see if a flag is present, or if a flag has been set in the command sequence for time-stamping of the data, or if the controller has been configured for time-stamping (Step 54). If the flag is not set to time-stamp the data it is written to storage media (Step 56).

If the flag is set to time-stamp the data it is passed to the time-stamping module (Step 58). The data is time-stamped (Step 60) and a digital signature effectively scaling the digital time-stamp to the data content, is applied (Step 62). A public key corresponding to this signature process can be placed on a network (Step 62a), e-mailed to a recipient of the data (Step 62b) or stored on media and mailed to a recipient of the data (Step 62c).

Alternatively, the public key can be recorded manually, not published at all, or published at any stage of the process.

The data timestamp and signature are then passed back to the controller (Step 64) and the time-stamping process is verified (Step 66). The data, time-stamp, and signature are then written to media (Step 68).

The coupling of the time-stamping features with a storage device ensures that data can always be securely written by this device and does not depend upon the application hosting server to provide secure data management. This is particularly useful in storage architectures which physically and
5 logically separate storage systems from application servers, e.g. storage area networks and network attached storage devices. All data written by the storage device can be content integrity checked and date/time of creation verified at a later date by decrypting and validation of the related signed time-stamp.

10

As can be seen from Figure 5, the present invention can reduce network traffic by removing the need to pass time-stamped data back across the network as it is time-stamped at the point at which it is stored.

15 Figure 5 shows a data originator 80 (e.g. computer, such as PC) connected to the Internet 81 via public telecommunications 82. Data to be time-stamped, signed and stored by a trusted clock data storage device is transmitted via public telecommunications 83 or 84 to a data storage device 85 or 86. In case of storage device 85, the trusted clock, signing capability,
20 and physical data store are all in one physical device, device 85, and the data is time-stamped signed and stored in device 85. In the case of device 86, the trusted clock and signing unit are in one physical box 87 and the memory is in another 88, or the memory may even be distributed memory 89 in a local network (not back out on the internet). This memory could be
25 disc or tape-based, or chip based. Of course, whilst the time-stamping and signing can be performed in the same "box", e.g. box 87, the signing could be in a different physical unit than the time-stamping, either in its own unit, or in the memory unit (still not requiring further access to the internet).

30 Data need only be passed to the time-stamping device and need not be passed back over the network once time-stamped for storage as the time-stamper and storage device (assembly, apparatus or system) are the same.

If the network is set up exclusively for the purpose of time-stamping network traffic can be halved. If it is a general purpose network the network traffic associated with time-stamping can still be significantly reduced.

5

Figure 6 shows a data storage device 90 having an interface I, a buffer 91, a trusted clock time-stamper 92, a controller 93, a signer 94, and a memory store 95. The controller 93 receives data from the buffer, decides what part of the data is to be time-stamped and sends that to the trusted clock 92 and receives back time-stamped data. The controller then sends the time-stamped data to the signer which signs it (creates a digest and encrypts the digest to create a signature). The signer then sends the signed time-stamped data back to the controller which sends it to memory 95 for storage.

15

In modified versions the signer could send the signed time-stamped data to the memory 95 without going through the controller. The clock 92 could send time-stamped data straight to the signer without going through the controller.

20

It will be appreciated that the controller may send all data to the clock for time-stamping, or just some data, e.g. selected types of data/selected parts of data. The time-stamper may stamp all data that it receives, or only some of the data that it receives. Data that is not time-stamped may or may not be recorded to memory.

25

Instead of the signing happening in the clock unit itself, it could occur externally of the clock unit, but still within the data storage device.

30

It will be appreciated that having a trusted clock attached to the data memory store provides the shortest path post-time-stamping/signing, which provides the least opportunity for attack on the integrity of the data and/or

timestamp, and the least opportunity for breakdowns or bottlenecks in external telecommunication systems to hinder the time-stamping and storage operation. Problems with congested networks hindering acquisition of a timestamp are similarly reduced if, once received by the data storage
5 system, the data does not have to go back out on an external network (e.g. the internet) for time-stamping and signing. Similarly, once time-stamped the data does not have to be subjected to Internet congestion/transmission problems before it is stored.

- 10 In some embodiments the trusted clock may be a device with a resonating crystal specifically intended for timekeeping. In other devices the clock may be a software clock, which may make use of the clock-speed of a processor chip. In either case, correction for drift of the clock may be possible, for example synchronisation with an external clock signal (e.g.
15 another trusted clock), possibly by wireless communication, possibly by wired (e.g. temporarily wired) connection.